

PCT


WORLD INTELLECTUAL PROPERTY ORGANIZATION
International Bureau



INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification 7 : G06F 9/38		A1	(11) International Publication Number: WO 00/43869
			(43) International Publication Date: 27 July 2000 (27.07.00)
(21) International Application Number: PCT/US00/01500			(81) Designated States: JP, European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE).
(22) International Filing Date: 24 January 2000 (24.01.00)			
(30) Priority Data: 09/236,081 25 January 1999 (25.01.99) US			
(71) Applicant: SUN MICROSYSTEMS, INC. [US/US]; 901 San Antonio Road, Palo Alto, CA 94303-4900 (US).			
(72) Inventor: TALCOTT, Adam, R.; 704 Villa Centre Way, San Jose, CA 95128 (US).			
(74) Agents: GARRETT, Arthur, S.; Finnegan, Henderson, Farabow, Garrett & Dunner, L.L.P., 1300 I Street, N.W., Washington, DC 20005-3315 (US) et al.			Published <i>With international search report.</i> <i>Before the expiration of the time limit for amending the claims and to be republished in the event of the receipt of amendments.</i>

(54) Title: METHODS AND APPARATUS FOR BRANCH PREDICTION USING HYBRID HISTORY WITH INDEX SHARING



(57) Abstract

A branch prediction scheme predicts whether a computer instruction will cause a branch to a non-sequential instruction. A prediction counter is selected by performing an exclusive or (XOR) operation between bits from an instruction address and a hybrid history. The hybrid history, in turn, is derived by concatenating bits from a global history register with bits from a local branch history table. The bits from the local branch history table are accessed by using bits from the instruction address.

FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AL	Albania	ES	Spain	LS	Lesotho	SI	Slovenia
AM	Armenia	FI	Finland	LT	Lithuania	SK	Slovakia
AT	Austria	FR	France	LU	Luxembourg	SN	Senegal
AU	Australia	GA	Gabon	LV	Latvia	SZ	Swaziland
AZ	Azerbaijan	GB	United Kingdom	MC	Monaco	TD	Chad
BA	Bosnia and Herzegovina	GE	Georgia	MD	Republic of Moldova	TG	Togo
BB	Barbados	GH	Ghana	MG	Madagascar	TJ	Tajikistan
BE	Belgium	GN	Guinea	MK	The former Yugoslav Republic of Macedonia	TM	Turkmenistan
BF	Burkina Faso	GR	Greece	ML	Mali	TR	Turkey
BG	Bulgaria	HU	Hungary	MN	Mongolia	TT	Trinidad and Tobago
BJ	Benin	IE	Ireland	MR	Mauritania	UA	Ukraine
BR	Brazil	IL	Israel	MW	Malawi	UG	Uganda
BY	Belarus	IS	Iceland	MX	Mexico	US	United States of America
CA	Canada	IT	Italy	NE	Niger	UZ	Uzbekistan
CF	Central African Republic	JP	Japan	NL	Netherlands	VN	Viet Nam
CG	Congo	KE	Kenya	NO	Norway	YU	Yugoslavia
CH	Switzerland	KG	Kyrgyzstan	NZ	New Zealand	ZW	Zimbabwe
CI	Côte d'Ivoire	KP	Democratic People's Republic of Korea	PL	Poland		
CM	Cameroon	KR	Republic of Korea	PT	Portugal		
CN	China	KZ	Kazakhstan	RO	Romania		
CU	Cuba	LC	Saint Lucia	RU	Russian Federation		
CZ	Czech Republic	LJ	Liechtenstein	SD	Sudan		
DE	Germany	LK	Sri Lanka	SE	Sweden		
DK	Denmark	LR	Liberia	SG	Singapore		
EE	Estonia						

METHODS AND APPARATUS FOR BRANCH PREDICTION USING HYBRID HISTORY WITH INDEX SHARING

BACKGROUND OF THE INVENTION

A. Field of the Invention

The invention generally relates to computer architecture, and, more particularly, to branch prediction.

B. Description of the Related Art

Modern high performance computer processors typically employ pipelining to increase performance. "Pipelining" refers to a processing technique in which multiple sequential instructions are executed in an overlapping manner. A general description of pipelining can be found in "Computer Organization & Design" by David A. Patterson and John L. Hennessy (2d ed. 1988, pp. 436-516).

Fig. 1 shows the timing of instruction processing in a conventional five-stage pipeline processor architecture. With such an architecture, the processor can simultaneously process different stages of up to five successive instructions. The five stages shown in Fig. 1 are: IF (instruction fetch), ID (instruction decode), EX (execute instruction), MEM (memory access), and WB (write back to register).

For example, at clock cycle 1, the processor fetches instruction I1. At clock cycle 2, the processor decodes instruction I1 and fetches instruction I2. In the same manner, the processor continues to process instructions as they are received; by clock cycle 5, the processor writes back the result of instruction I1, accesses memory for instruction I2, executes instruction I3, decodes instruction I4, and fetches instruction I5. In contrast, a non-pipelined architecture would complete processing of an entire instruction (e.g., instruction I1) before beginning to process the next instruction (e.g., instruction I2).

When program flow is perfectly sequential, a pipelined architecture can achieve significant performance advantages over non-pipelined architecture. In actual programs, however, approximately twenty percent of program instructions are branches. Branch instructions cause a program to deviate from a sequential flow. Consequently, the instruction to be executed (the target of the branch) may not be the next instruction in the fetch sequence.

A processor may recognize that an instruction is a branch instruction in the IF stage (the first stage of the five-stage pipeline). For conditional branch instructions, however, the processor typically cannot determine whether the branch should be taken until it reaches the EX stage (the third stage of the five-stage pipeline). By this time, the processor has already fetched and begun processing the next two instructions. The processing of those two instructions is wasted and inefficient if the branch instruction redirects program flow to another location.

Referring to Fig. 1, if instruction I1 is a conditional branch instruction that redirects flow to instruction I6, the processor does not recognize this until clock cycle 3 (EX), when the processor is executing instruction I1. By this time, the processor has already fetched instruction I2 during clock cycle 2, and decoded instruction I2 and fetched instruction I3 during clock cycle 3. This processing of instructions I2 and I3 is wasted, however, because branch instruction I1 causes flow to skip to instruction I6, with no further processing of instructions I2 or I3. Moreover, the branching causes a stall in the pipeline while the correct instruction (I6) is fetched. These inefficiencies caused by branches become exacerbated when deeper pipelines or superscalar processors are used because it takes longer to resolve a branch.

One approach to solving this problem, called branch prediction, involves making accurate, educated determinations about whether an instruction will result in a branch to another location. Branch prediction is premised on the assumption that, under similar circumstances, the outcome of a conditional branch will likely be the same as prior outcomes. Because branch prediction can be implemented in the IF stage of processing, there is no wasted instruction processing if the result of the conditional branch is always predicted correctly.

Conventional branch prediction techniques include correlation-based schemes and global branch history with index sharing ("gshare"). Although these techniques are somewhat effective, the frequency of erroneous prediction using these techniques may be unacceptable. There remains, therefore, a need for a branch prediction scheme that reduces the frequency of erroneous prediction.

SUMMARY OF THE INVENTION

In accordance with the invention, as embodied and broadly described herein, a method of predicting whether a branch will be taken involves reading bits from a local history table and concatenating them with bits from a global history register. The result of the concatenation is combined with bits from the instruction address by performing an exclusive or operation. The result of the exclusive or operation is used to read a branch prediction table.

In accordance with the invention, an apparatus for predicting whether a branch will be taken comprises a local history table and a global history register. The local history table and the global history table are connected to inputs of a concatenating circuit. The output of the concatenating circuit is connected to one input of an exclusive or circuit, with an instruction address source being connected to another input. The output of the exclusive or circuit is connected to an input of a branch prediction table.

It is to be understood that both the foregoing general description and following detailed description are intended only to exemplify and explain the invention as claimed.

BRIEF DESCRIPTION OF THE DRAWINGS

The accompanying drawings, which are incorporated in and constitute a part of this specification, illustrate embodiments of the invention and, together with the description, serve to explain the advantages and principles of the invention. In the drawings,

FIG. 1 shows the timing of instruction processing in a conventional five-stage pipeline processor architecture;

FIG. 2 is a block diagram depicting a first system for branch prediction, consistent with the invention;

FIG. 3 is a block diagram depicting a second system for branch prediction, consistent with the invention;

FIG. 4 is a flow diagram of steps performed to predict whether a branch will be taken, consistent with the invention; and

FIG. 5 is another flow diagram of steps performed to predict whether a branch will be taken, consistent with the invention.

DETAILED DESCRIPTION

The invention will now be described in reference to the accompanying drawings. The same reference numbers may be used throughout the drawings and the following description to refer to the same or like parts.

A. Overview

Methods and apparatus consistent with the invention predict whether an instruction will cause a branch to a non-sequential instruction. This is achieved by incorporating features of the correlation-based and gshare schemes to obtain a scheme consistent with the invention. In particular, a prediction counter is selected by performing an exclusive or (XOR) operation between (i) a specified number of bits from an instruction address and (ii) a hybrid history. The hybrid history, in turn, is derived by concatenating (i) a specified number of bits from a global history register with (ii) a specified number of bits from a local branch history table. The bits from the local branch history table are accessed by using a specified number of bits from the instruction address.

B. Architecture

Fig. 2 is a block diagram depicting one system of branch prediction, consistent with the invention. In a preferred embodiment, system 200 includes the following components: local history table 220, concatenator 230, global history register 240, XOR 250, and branch prediction table 260. System 200 may be controlled and accessed by an instruction fetch unit ("IFU") 290.

Local history table 220 is connected to concatenator 230 via a data path that is preferably l bits wide. Global history register 240 is also connected to concatenator 230 via a data path that is preferably g bits wide. Concatenator 230 is connected to XOR 250 via a data path that is preferably $l+g$ bits wide. XOR 250 is connected to branch prediction table 260 via a data path that is preferably $l+g$ bits wide.

Local history table 220 is a device storing local history data and preferably comprises 2^a shift registers, each having at least l bits. Alternatively, local history data may be stored in any type of memory, such as a single register, multiple registers, or random access memory. Each register stores the l most recent conditional outcomes for a set of instruction addresses that each have the same a address bits in common. When an

instruction in this set results in a branch being taken, a value of 1 is shifted into the corresponding register. In contrast, a value of 0 is shifted into the corresponding register if a branch is not taken. Data that corresponds to branch history on a local level is hereinafter called "local branch history data."

Global history register 240 preferably comprises a shift register having at least g bits. These bits represent the g most recent outcomes for any branch instruction, conditional or not, and regardless of its address. When a branch is taken, a value of 1 is shifted into global history register 240. In contrast, a value of 0 is shifted into global history register 240 when a branch is not taken. Data that corresponds to branch history at a global level is hereinafter called "global branch history data."

Concatenator 230 is a device that receives g bits from global history register 240 and l bits from local history table 220, and concatenates them together to form an output having $l+g$ bits. XOR 250 is a device that receives two inputs each having $l+g$ bits, performs an exclusive or (XOR) operation between the two inputs on a bit-by-bit basis, and creates an output having $l+g$ bits.

Branch prediction table 260 is a device storing branch prediction data and may be implemented using a plurality of n -bit saturating counters. Each of these counters stores data representing whether a branch was taken under a particular circumstance. A circumstance may be defined by the input to branch prediction table 260, which in system 200 may be based on the values of the instruction address, the global history register, and the local history register. For a particular circumstance, if a branch is taken, the value of the corresponding counter is incremented; if the counter value is already at its maximum value, it remains there. If, on the other hand, a branch is not taken, the value of the counter corresponding to that circumstance is decremented; if the counter value is already at its minimum value, it remains unchanged.

These n -bit saturating counters are the basis for the branch prediction decision. For a particular circumstance, if a branch was previously taken (indicated by the counter having a designated value), system 200 predicts that the branch will be taken again. In a preferred embodiment, system 200 predicts that the branch will be taken if the most significant bit of the n -bit counter equals "1".

In one embodiment, $a=14$, $l=4$, $g=12$, and $n=2$. Other values for these variables can also be used. For example, in another embodiment, $a=14$, $l=2$, $g=14$, and $n=2$.

Fig. 3 is block diagram depicting another system of branch prediction, consistent with the present invention. In a preferred embodiment, system 300 includes the following components: local history table 220, global history register 240, branch prediction table 260, XOR 310, XOR 320, and column decode multiplexer 330. System 300 may be controlled and accessed by IFU 290.

Global history register 240 is connected to XOR 310 via a data path that is preferably g bits wide. XOR 310 is connected to branch prediction table 260 via a data path that is preferably g bits wide. Branch prediction table 260 is connected to multiplexer 330 via 2^l data paths that are preferably n bits wide. The select line of multiplexer 330 is connected to XOR 320 via a data path that is l bits wide. XOR 320 is connected to local history table 320 via a data path that is l bits wide.

XOR 310 is a device that receives two inputs that are g bits wide, performs an XOR operation on them on a bit-by-bit basis, and generates an output that is also g bits wide. Similarly, XOR 320 is a device that receives two inputs that are l bits wide, performs an XOR operation on them on a bit-by-bit basis, and generates an output that is also l bits wide. Multiplexer 330 receives 2^l inputs that are n -bits wide. In response to a control signal from XOR 320, multiplexer 330 passes along one of the 2^l inputs. The system shown in Fig. 3 may use the same values for variables a , l , g , and n as stated above in reference to Fig. 2.

Systems 200 and 300 described in reference to Figs. 2 and 3 may be implemented as an integrated circuit as part of one or more computer processors. Alternatively, systems 200 and 300 may be implemented in discrete logic components or software and may be implemented separate from a processor.

C. Architectural Operation

Fig. 4 is a flow diagram of a process consistent with the invention, and is described with reference to system 200 shown in Fig. 2.

The process begins with IFU 290 reading local history table 220 (step 410). Specifically, IFU uses bits $a+1:2$ of the instruction address to access local history data

from local history table 220. As used herein, the terminology " $m:n$ " denotes bits m through n , inclusive. In a preferred embodiment, these bits correspond to the a least significant bits of the instruction address excluding the last two bits. The last two bits are preferably excluded because they are typically zero in a processor that employs byte addressing and 32-bit instructions. By accessing local history table 220, IFU 290 causes it to generate an output that is at least l bits wide.

Those of skill in the art will recognize that the invention is not limited to using a portion of the instruction address to access the local history table. For example, the local history table could instead be accessed based upon an address that corresponds to a group of instructions.

Concatenator 230 concatenates the l -bit output from local history table 220 with g bits from global history register 240 (step 420). The output of concatenator 230 may be referred to as either concatenated history data or a hybrid history. XOR 250 performs an XOR operation between the $l+g$ bits output by concatenator 230 and $l+g+1:2$ bits from the instruction address (step 430). These $l+g+1:2$ bits correspond to the $l+g$ least significant bits of the instruction address, preferably excluding the last two bits.

IFU 290 uses the $l+g$ bits resulting from the XOR operation to read branch prediction table 260 (step 440). In response, branch prediction table 260 generates an output that is n bits wide. IFU 290 then interprets this n -bit output to predict whether a branch will occur (step 450). Specifically, if the n -bit counter indicates that a branch was taken previously under similar circumstances, then a prediction is made that the branch will again be taken.

Fig. 5 is a flow diagram of another process consistent with the invention. The process shown in Fig. 5 is described with reference to system 300, shown in Fig. 3.

The process begins with XOR 310 performing an XOR operation between g bits from global history register 240 and $l+g-3:l+2$ bits from the instruction address (step 510). These $l+g-3:l+2$ bits correspond to the same portion of the instruction address that is used in the XOR operation with the global history portion of the hybrid history in FIG. 2. IFU 290 uses this g -bit output from exclusive or 310 as an input to read branch

prediction table 260 (step 520). In response, branch prediction table 260 generates 2^l outputs that are each n -bits wide. These outputs are fed as inputs to multiplexer 330.

Concurrent with the read of branch prediction table 260, the system performs a read of local history table 220 (step 530). Specifically, IFU 290 reads local history table 220 via bits $a+1:2$ of the instruction address; the last two bits are preferably excluded as discussed above in connection with the process shown in Fig. 3. In response to this read operation, local history table 220 generates an output that is l bits wide. XOR 320 performs an XOR operation between this l -bit output and bits $l+1:2$ from the instruction address (step 540). Again, the last two bits are preferably excluded. This creates an output from exclusive or 320 that is l -bits wide.

IFU 290 uses the l -bit output from exclusive or 320 as a "select" input to multiplexer 330. In response, multiplexer 330 generates an n -bit output equivalent to one of the 2^l outputs of branch prediction table 260 (step 550). IFU 290 then interprets this n -bit output to predict whether or not a branch will occur (step 560). Specifically, if the most significant bit of n -bit counter indicates that a branch was taken previously under similar circumstances, then a prediction is made that the branch will again be taken. In a preferred embodiment, if the most significant bit of n -bit counter equals "1", then a prediction is made that the branch will again be taken.

D. Conclusion

As described in detail above, methods and apparatus consistent with the present invention predict whether a branch will be taken. The foregoing description of an implementation of the invention has been presented for purposes of illustration and description. It is not exhaustive and does not limit the invention to the precise form disclosed. Modifications and variations are possible in light of the above teachings or may be acquired from practicing the invention. For example, the data paths between the various components may be in the form of integrated circuit connections, wires, or fiber optics, to name a few. Similarly, although the description above is based on a processor that employs byte addressing and 32-bit instructions, a similar approach could be employed with other addressing schemes. Moreover, the description above is based on a single-processor pipeline architecture, but the invention may be used in a multiple

processor environment and non-pipeline processor environments. Furthermore, although the description above employs an XOR function, other hashing functions could be used consistent with the invention. The scope of the invention is defined by the claims and their equivalents.

What is claimed is:

1. A method of predicting whether processing of an instruction is to result in branching of program flow, comprising:

accessing local branch history data;

accessing global branch history data;

concatenating the accessed local branch history data with the accessed global branch history data to form concatenated history data;

performing an exclusive or operation between the concatenated history data and bits from the address of the instruction to form an index to branch prediction data; and

generating branch prediction data based upon the index.

2. The method of claim 1, further comprising interpreting the branch prediction data to determine whether a branch will be taken.

3. The method of claim 1, further comprising updating the global branch history data based upon whether a branch is taken.

4. The method of claim 1, further comprising updating the local branch history data based upon whether a branch is taken.

5. A method of predicting whether processing of an instruction is to result in branching of program flow, comprising:

accessing global branch history data;

performing a first exclusive or operation between the accessed global branch history data and bits from the address of the instruction to form index data;

accessing local branch history data;

performing a second exclusive or operation between the accessed local branch history data and bits from the address of the instruction to form select data;

accessing branch prediction data based upon the index data; and

selecting at least one of the branch prediction data based upon the select data.

6. A method of predicting whether processing of an instruction is to result in branching of program flow, comprising:

reading local branch history data from a local history table;

reading global branch history data from a global history register;

concatenating the read local branch history data with the read global branch history data to form concatenated history data;

performing an exclusive or operation between the concatenated history data and bits from the address of the instruction to form an index to branch prediction data; and

reading branch prediction data from a branch prediction table based upon the index.

7. A method of predicting whether processing of an instruction is to result in branching of program flow, comprising:

reading global branch history data from a global history register;

performing a first exclusive or operation between the read global branch history data and bits from the address of the instruction to form index data;

reading local branch history data from a local history table;

performing a second exclusive or operation between the read local branch history data and bits from the address of the instruction to form select data;

reading branch prediction data from a branch prediction table based upon the index data; and

selecting at least one of the branch prediction data based upon the select data.

8. A method of predicting whether processing of an instruction is to result in branching of program flow, comprising:

providing a portion of an instruction address as an input to a local history table;

providing an output of a local history table as a first input to a concatenating circuit;

providing an output of a global history register as a second input to the concatenating circuit;

providing an output from the concatenating circuit as a first input to an exclusive or circuit;

providing a portion of the instruction address as a second input to the exclusive or circuit; and

providing an output from the exclusive or circuit as an input to a branch prediction table.

9. A method of predicting whether processing of an instruction is to result in branching of program flow, comprising:

providing an output of a global history register as a first input to a first exclusive or circuit;

providing a portion of an instruction address as a second input to the first exclusive or circuit;

providing an output from the first exclusive or circuit as an input to a branch prediction table;

providing a portion of an instruction address as an input to a local history table;

providing an output of the local history table as a first input to a second exclusive or circuit;

providing a portion of an instruction address as a second input to a second exclusive or circuit; and

providing an output from the second exclusive or circuit as an input to a select circuit.

10. An apparatus for predicting whether processing of an instruction is to result in branching of program flow, comprising:

means for accessing local branch history data;

means for accessing global branch history data;

means for concatenating the accessed local branch history data with the accessed global branch history data to form concatenated history data;

means for performing an exclusive or operation between the concatenated history data and bits from the address of the instruction to form an index to branch prediction data; and

means for generating branch prediction data based upon the index.

11. An apparatus for predicting whether processing of an instruction is to result in branching of program flow, comprising:

means for accessing global branch history data;

means for performing a first exclusive or operation between the accessed global branch history data and bits from the address of the instruction to form index data;

means for accessing local branch history data;

means for performing a second exclusive or operation between the accessed local branch history data and bits from the address of the instruction to form select data;

means for accessing branch prediction data based upon the index data; and

means for selecting at least one of the branch prediction data based upon the select data.

12. An apparatus for predicting whether processing of an instruction is to result in branching of program flow, comprising:

means for reading local branch history data from a local history table;

means for reading global branch history data from a global history register;

means for concatenating the read local branch history data with the read global branch history data to form concatenated history data;

means for performing an exclusive or operation between the concatenated history data and bits from the address of the instruction to form an index to branch prediction data; and

means for reading branch prediction data from a branch prediction table based upon the index.

13. An apparatus for predicting whether processing of an instruction is to result in branching of program flow, comprising:

means for reading global branch history data from a global history register;

means for performing a first exclusive or operation between the read global branch history data and bits from the address of the instruction to form index data;

means for reading local branch history data from a local history table;

means for performing a second exclusive or operation between the read local branch history data and bits from the address of the instruction to form select data;

means for reading branch prediction data from a branch prediction table based upon the index data; and

means for selecting at least one of the branch prediction data based upon the select data.

14. An apparatus for predicting whether processing of an instruction is to result in branching of program flow, comprising:

means for providing a portion of an instruction address as an input to a local history table;

means for providing an output of a local history table as a first input to a concatenating circuit;

means for providing an output of a global history register as a second input to the concatenating circuit;

means for providing an output from the concatenating circuit as a first input to an exclusive or circuit;

means for providing a portion of an instruction address as a second input to the exclusive or circuit; and

means for providing an output from the exclusive or circuit as an input to a branch prediction table.

15. An apparatus for predicting whether processing of an instruction is to result in branching of program flow, comprising:

means for providing an output of a global history register as a first input to a first exclusive or circuit;

means for providing a portion of an instruction address as a second input to the first exclusive or circuit;

means for providing an output from the first exclusive or circuit as an input to a branch prediction table;

means for providing a portion of an instruction address as an input to a local history table;

means for providing an output of the local history table as a first input to a second exclusive or circuit;

means for providing a portion of an instruction address as a second input to a second exclusive or circuit; and

means for providing an output from the second exclusive or circuit as an input to a select circuit.

16. An apparatus for predicting whether processing of an instruction is to result in branching of program flow, comprising:

a first memory storing local branch history data;

a second memory storing global branch history data;

a third memory storing branch prediction data;

a concatenating device having first and second inputs connected to the first memory and the second memory, respectively, and an output; and

a XOR device having a first input connected to the output of the concatenating device, a second input receiving at least a portion of an address of the instruction, and an output connected to the third memory.

17. An apparatus for predicting whether processing of an instruction is to result in branching of program flow, comprising:

a local branch history table;

a global branch history register;

a branch prediction table;

a concatenating device having first and second inputs connected to the local branch history table and the global branch history register, respectively, and an output; and

a XOR device having a first input connected to the output of the concatenating device, a second input receiving at least a portion of an address of the instruction, and an output connected to the branch prediction table.

18. An apparatus for predicting whether processing of an instruction is to result in branching of program flow, comprising:

a first memory storing local branch history data;

a second memory storing global branch history data;

a third memory storing branch prediction data;

a first XOR device having a first input connected to the second memory, a second input receiving at least a portion of an address of the instruction, and an output connected to the third memory; and

a second XOR device having a first input connected to the first memory, a second input receiving at least a portion of an address of the instruction, and an output connected to a multiplexer.

19. An apparatus for predicting whether processing of an instruction is to result in branching of program flow, comprising:

a local branch history table;

a global branch history register;

a branch prediction table;

a first XOR device having a first input connected to the global branch history register, a second input receiving at least a portion of an address of the instruction, and an output connected to the branch prediction table; and

a second XOR device having a first input connected to the local branch history table, a second input receiving at least a portion of an address of the instruction, and an output connected to a select line of a multiplexer.

20. A system that predicts whether processing of an instruction is to result in branching of program flow, comprising:

a processor for executing instructions;

a first memory storing local branch history data and connected to the processor;

a second memory storing global branch history data;

a third memory storing branch prediction data and connected to the processor;

a concatenating device having first and second inputs connected to the first memory and the second memory, respectively, and an output; and

a XOR device having a first input connected to the output of the concatenating device, a second input receiving at least a portion of an address of the instruction, and an output connected to the third memory.

21. The system according to claim 20, wherein the processor is configured to execute instructions in a pipeline.

22. The system according to claim 21, wherein the processor comprises:

an instruction fetch unit for fetching instructions;

an instruction decode unit for decoding fetched instructions;

an execution unit for executing the decoded instructions;

a memory access unit for accessing data from a memory; and

a write back unit to write data to a memory.

23. A system that predicts whether processing of an instruction is to result in branching of program flow, comprising:

an instruction fetch unit for fetching instructions to be processed;

a first memory storing local branch history data and connected to the instruction fetch unit;

a second memory storing global branch history data;

a third memory storing branch prediction data and connected to the instruction fetch unit;

a concatenating device having first and second inputs connected to the first memory and the second memory, respectively, and an output; and

a XOR device having a first input connected to the output of the concatenating device, a second input receiving at least a portion of an address of the instruction, and an output connected to the third memory.

Clock Cycle		1	2	3	4	5	6	7	8
Instruction	I1	IF	IF	EX	MEM	WB			
	I2		IF	ID	EX	MEM	WB		
	I3			IF	ID	EX	MEM	WB	
	I4				IF	ID	EX	MEM	WB
	I5					IF	ID	EX	MEM
	I6						IF	ID	EX

FIG. 1


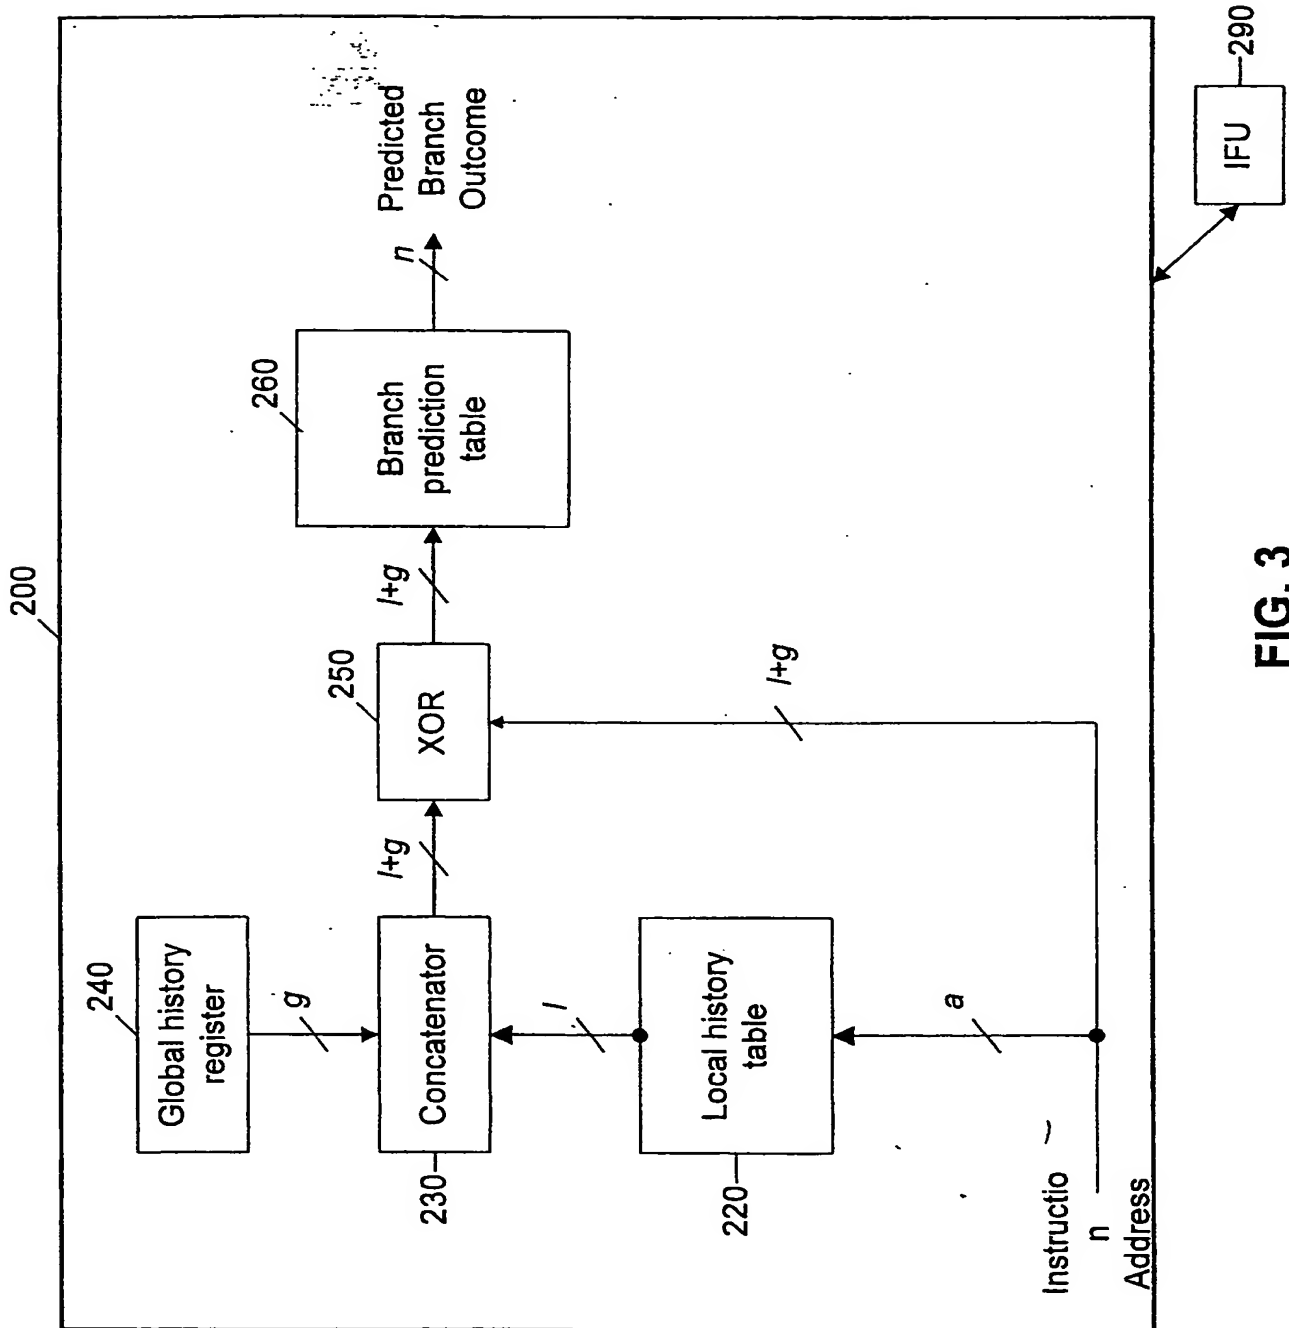


FIG. 2



4/5

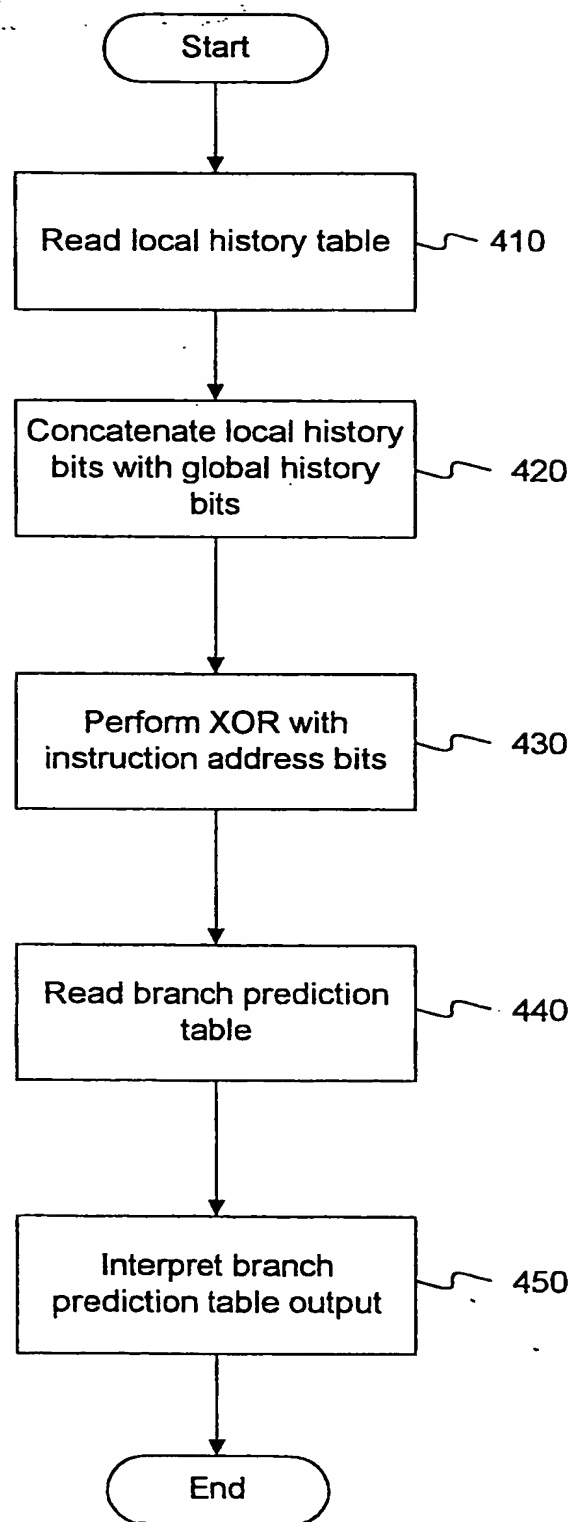


FIG. 4

5/5

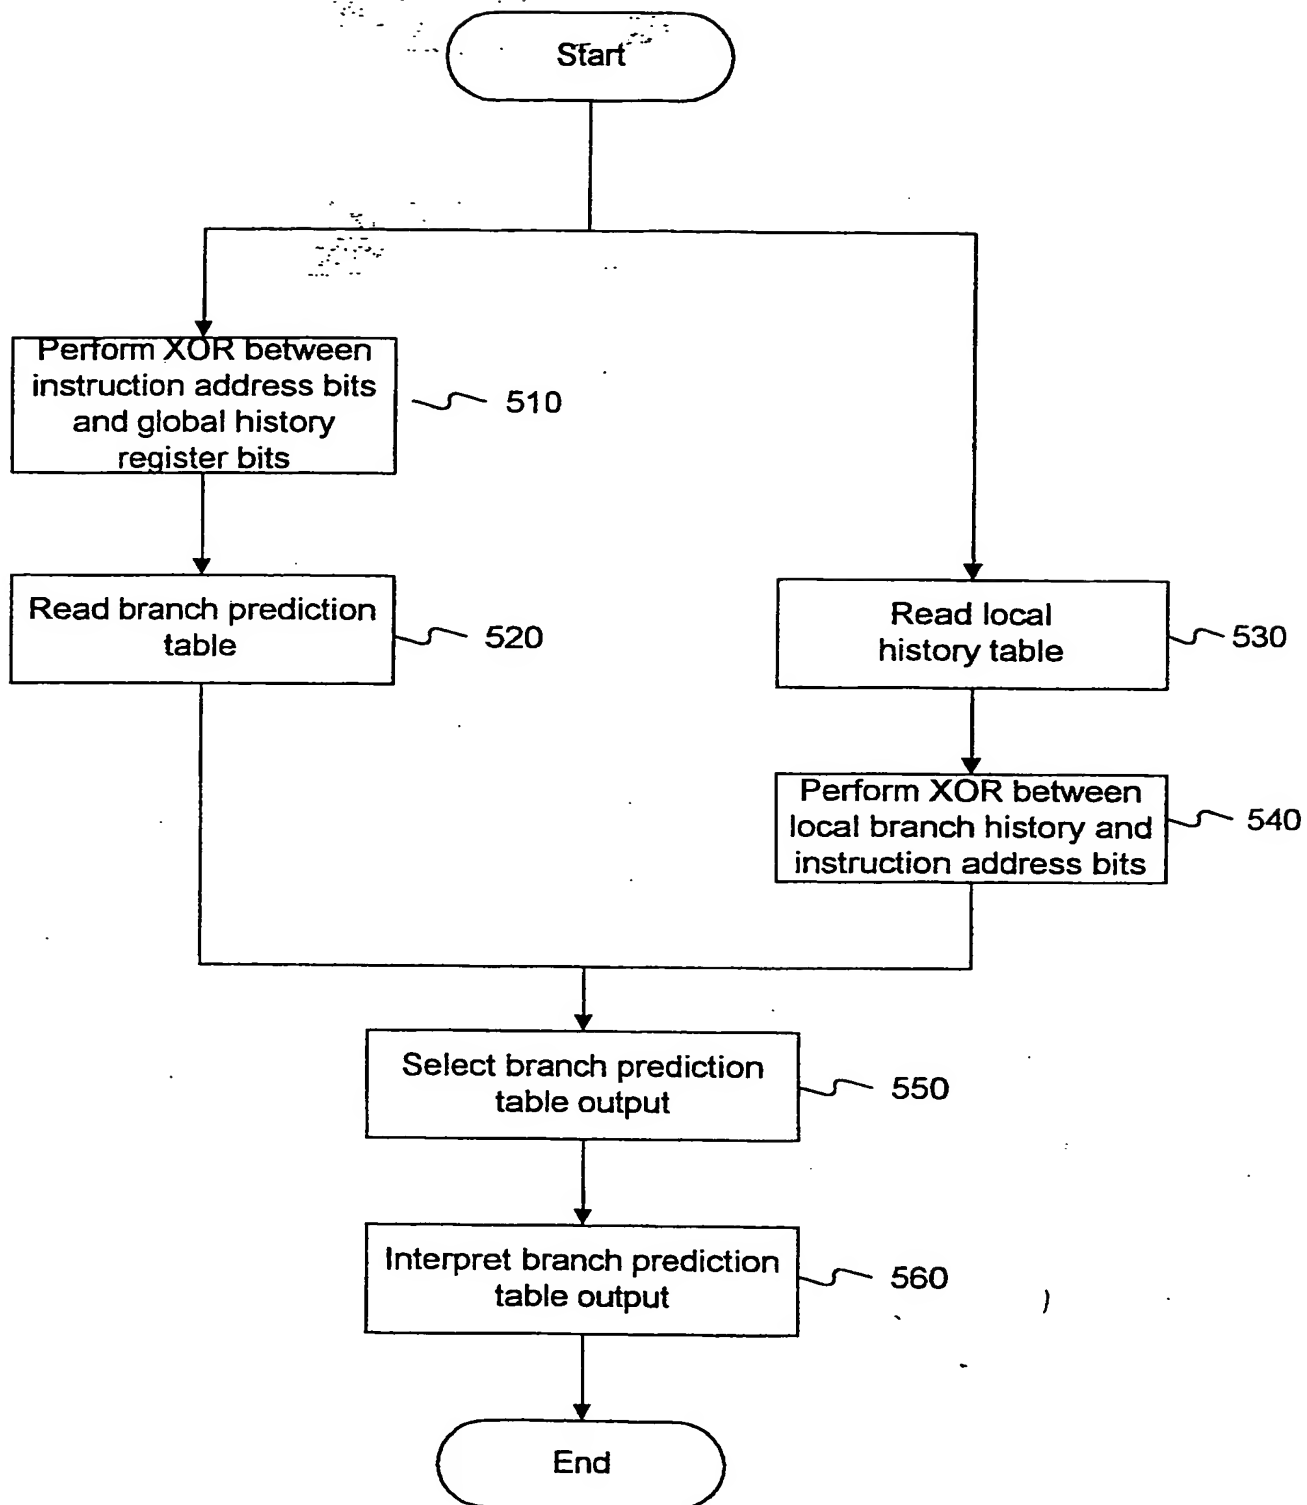


FIG. 5

INTERNATIONAL SEARCH REPORT

Int'l Application No
PCT/US 00/01500

A. CLASSIFICATION OF SUBJECT MATTER

IPC 7 606F9/38

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

IPC 7 606F

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	EP 0 731 408 A (INT COMPUTERS LTD) 11 September 1996 (1996-09-11) the whole document	1,5-20, 23
A	WO 97 30389 A (HAL COMPUTER SYSTEMS INC) 21 August 1997 (1997-08-21) page 7, line 22 -page 9, line 24 -/-	1,5-20, 23

☒ Further documents are listed in the continuation of box C.

☒ Patent family members are listed in annex.

* Special categories of cited documents :

"A" document defining the general state of the art which is not considered to be of particular relevance

"E" earlier document but published on or after the international filing date

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.

"&" document member of the same patent family

Date of the actual completion of the international search

9 June 2000

Date of mailing of the international search report

16/06/2000

Name and mailing address of the ISA

European Patent Office, P.B. 5818 Patentaan 2
NL - 2280 HV Rijswijk
Tel. (+31-70) 340-2040, Tx. 31 651 epo nl,
Fax: (+31-70) 340-3016

Authorized officer

Moraiti, M

INTERNATIONAL SEARCH REPORT

Inte nal Application No
PCT/US 00/01500

C.(Continuation) DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	<p>FAGIN B ET AL: "THE PERFORMANCE OF COUNTER- AND CORRELATION-BASED SCHEMES FOR BRANCH TARGET BUFFERS" IEEE TRANSACTIONS ON COMPUTERS,US,IEEE INC. NEW YORK, vol. 44, no. 12, 1 December 1995 (1995-12-01), pages 1383-1393, XP000541980 ISSN: 0018-9340 page 4, column 2, paragraph 4 -page 5, column 1, paragraph 1</p>	1,5,6
A	<p>TALCOTT A R ET AL: "THE IMPACT OF UNRESOLVED BRANCHES ON BRANCH PREDICTION SCHEME PERFORMANCE" PROCEEDINGS OF THE ANNUAL INTERNATIONAL SYMPOSIUM ON COMPUTER ARCHITECTURE,US,LOS ALAMITOS, IEEE COMP. SOC. PRESS, vol. SYMP. 21, 1994, pages 12-21, XP000480422 ISBN: 0-8186-5510-0 page 2, column 2, paragraph 2 -page 3, column 1, paragraph 1</p>	
A	<p>PO-YUNG CHANG ET AL: "ALTERNATIVE IMPLEMENTATIONS OF HYBRD BRANCH PREDICTORS" PROCEEDINGS OF THE ANNUAL INTERNATIONAL SYMPOSIUM ON MICROARCHITECTURE,US,LOS ALAMITOS, IEEE COMP. SOC. PRESS, vol. SYMP. 28, 1995, pages 252-257, XP000585367 ISBN: 0-8186-7349-4 page 4, column 1, paragraph 3 -column 2, paragraph 4</p>	1,5,6

INTERNATIONAL SEARCH REPORT

Information on patent family members

International Application No

PCT/US 00/01500

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
EP 0731408 A	11-09-1996	AU 700454 B	07-01-1999
		AU 4791796 A	19-09-1996
		DE 69601742 D	22-04-1999
		DE 69601742 T	14-10-1999
		US 5721894 A	24-02-1998
		ZA 9601534 A	27-08-1996
WO 9730389 A	21-08-1997	US 5742805 A	21-04-1998
		EP 0823085 A	11-02-1998
		JP 11504458 T	20-04-1999
		US 6055629 A	25-04-2000
		US 5896529 A	20-04-1999